



US010366116B2

(12) **United States Patent**
Byron et al.

(10) **Patent No.:** **US 10,366,116 B2**
(45) **Date of Patent:** ***Jul. 30, 2019**

(54) **DISCREPANCY CURATOR FOR DOCUMENTS IN A CORPUS OF A COGNITIVE COMPUTING SYSTEM**

(58) **Field of Classification Search**
CPC G06F 17/3071; G06F 17/30696; G06F 17/30699; G06F 17/2725; G06N 99/005
See application file for complete search history.

(71) Applicant: **International Business Machines Corporation**, Armonk, NY (US)

(56) **References Cited**

(72) Inventors: **Donna K. Byron**, Petersham, MA (US);
Elie Feirouz, Lexington, MA (US);
Ashok Kumar, Bedford, MA (US);
William G. O’Keeffe, Cambridge, MA (US)

U.S. PATENT DOCUMENTS

7,313,515 B2 12/2007 Crouch et al.
7,461,047 B2 12/2008 Masuichi et al.
(Continued)

(73) Assignee: **International Business Machines Corporation**, Armonk, NY (US)

FOREIGN PATENT DOCUMENTS

WO 2013125286 A1 8/2013

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 41 days.

OTHER PUBLICATIONS

Source Forge; “The Lemur Project”, retrieved on Feb. 7, 2017 from [https://sourceforge.net/p/lemur/wiki/Indexer File Formats](https://sourceforge.net/p/lemur/wiki/Indexer+File+Formats).

(Continued)

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **15/800,851**

Primary Examiner — Fariba Sirjani

(22) Filed: **Nov. 1, 2017**

(74) *Attorney, Agent, or Firm* — Robert H. Frantz;
William J. Stock

(65) **Prior Publication Data**

US 2018/0067950 A1 Mar. 8, 2018

(57) **ABSTRACT**

Curation of a corpus of a cognitive computing system is performed interactively by reporting on user interface device to a user a parse tree illustration of discrepancies and corresponding assigned confidence factors detected between at least a portion of a first document and a second or more documents in the corpus. Responsive to a user selection of an illustrated discrepancy in the parse tree, a drill-down dialog is prepared and displayed which shows at least a text string for the portion of the first document and at least one conflicting text string from the second or more documents, and which provides at least one user-selectable administrative action option for handling the detected discrepancy. Responsive to receipt of user selection of an administrative action option, the computing system performs the action to handle the detected discrepancy.

Related U.S. Application Data

(63) Continuation of application No. 14/993,538, filed on Jan. 12, 2016, now Pat. No. 9,842,161.


(51) **Int. Cl.**
G06F 16/35 (2019.01)
G06F 17/27 (2006.01)

(Continued)

(52) **U.S. Cl.**
CPC **G06F 16/355** (2019.01); **G06F 16/335** (2019.01); **G06F 16/338** (2019.01); **G06F 17/2725** (2013.01); **G06N 20/00** (2019.01)

18 Claims, 9 Drawing Sheets

800



View conflicts and take appropriate course of action to resolve them 801

Check GTT to ensure the mapped answer is still the best answer, add additional answers etc... 802

Add Ground Truth and train if appropriate 803

804

(51)	Int. Cl.		2017/0199929 A1 *	7/2017	Byron	G06F 16/334
	G06N 20/00	(2019.01)	2018/0001470 A1 *	1/2018	Rusu	B25J 9/163
	G06F 16/335	(2019.01)	2018/0260860 A1 *	9/2018	Devanathan	G06F 17/2785
	G06F 16/338	(2019.01)				

OTHER PUBLICATIONS

(56) References Cited

U.S. PATENT DOCUMENTS

8,370,275	B2	2/2013	Bhattacharya et al.	
8,560,300	B2	10/2013	Ferrucci et al.	
9,015,098	B1	4/2015	Crosley	
9,842,161	B2 *	12/2017	Byron	G06N 20/00
10,146,858	B2 *	12/2018	Byron	G06F 16/35
10,176,250	B2 *	1/2019	Byron	G06F 16/334
2002/0127521	A1	9/2002	Fegan	
2013/0138425	A1 *	5/2013	Luke	G06F 17/2705
				704/9
2014/0006557	A1	1/2014	Lindgren et al.	
2014/0046947	A1 *	2/2014	Jenkins	G06F 17/2881
				707/740
2014/0120513	A1	5/2014	Jenkins et al.	
2014/0172139	A1	6/2014	Clark et al.	
2014/0172756	A1	6/2014	Clark et al.	
2014/0172878	A1	6/2014	Clark et al.	
2014/0172880	A1	6/2014	Clark et al.	
2014/0172882	A1	6/2014	Clark et al.	
2014/0172883	A1	6/2014	Clark et al.	
2015/0026106	A1	1/2015	Oh et al.	
2015/0142418	A1 *	5/2015	Byron	G06F 17/2785
				704/9
2015/0286977	A1 *	10/2015	Schneeman	G06Q 10/06393
				705/7.39
2015/0324350	A1 *	11/2015	Bufe	G06F 17/28
				704/9
2016/0148114	A1 *	5/2016	Allen	G16H 10/60
				706/11
2017/0177675	A1 *	6/2017	Beller	G06F 16/90332
2017/0193174	A1 *	7/2017	Allen	G06F 17/2765

Stack Overflow; "Where does one find the TREC Document format specification?", retrieved on Feb. 7, 2017 from <http://www.stackoverflow.com>.

Stack Overflow; "What is the TREC format?", retrieved on Feb. 7, 2017 from <http://www.stackoverflow.com>.

Wikipedia; "Text Retrieval Conference", retrieved on Feb. 7, 2017 from <https://www.wikipedia.com>.

Yuan, Michael J.; "Watson and Healthcare", Apr. 12, 2011, retrieved on Dec. 3, 2015 from <http://www.ibm.com/developerworks/library/os-ind-watson/>.

High, Rob; "The Era of Cognitive Systems: An Inside Look at IBM Watson and How it Works", 2012, retrieved on Dec. 3, 2015 from www.redbooks.ibm.com/redpapers/pdfs/redp4955.pdf.

Temperley, Davy, et al.; "Link Grammar", retrieved on Nov. 30, 2015 from <http://www.link.cs.cmu.edu/link/>.

Byron, Donna K., et al.; "Utilizing a Dialectical Model in a Question Answering System"; U.S. Appl. No. 14/722,862, filed May 27, 2015.

Ferrucci, D. A., et al.; "This is Watson", IBM Journal of Research and Development, vol. 56, No. 3/4, May/Jul. 2012.

Ferrucci, David, et al.; "The AI Behind Watson—The Technical Article"; 2010, retrieved on Nov. 12, 2015 from <http://www.aaai.org/Magazine/Watson/watson.php>.

IBM; "Watson—A System Designed for Answers", Feb. 2011.

Kelley III, Joe. E.; "Computing, cognition and the future of knowing", Copyright 2015 by IBM Corporation.

Budnik, Tom, et al.; "High Throughput Computing on IBM's Blue Gene/P".

U.S. Appl. No. 14/993,538, filed Jan. 10, 2016, US20170199882A1.

U.S. Appl. No. 14/722,862, filed May 27, 2015, US20160350279A1.

U.S. Appl. No. 14/966,568, filed Dec. 11, 2015, US20170169017A1.

* cited by examiner




Fig. 1




Fig. 2




Fig. 3




Fig. 4




Fig. 5




Fig. 6




Fig. 7




Fig. 8

900




Fig. 9

1

DISCREPANCY CURATOR FOR DOCUMENTS IN A CORPUS OF A COGNITIVE COMPUTING SYSTEM

This is a continuation application of U.S. patent application Ser. No. 14/993,538, filed on Jan. 12, 2016, by Donna K. Byron, et al., now under Notice of Allowance. This invention relates generally to methods to align and reconcile documents in a cognitive computing system corpus of for which differences and contradictions have been found amongst other documents and information elements in the corpus.

INCORPORATION BY REFERENCE

U.S. patent application Ser. No. 14/722,862, filed on May 27, 2015, by Donna K. Byron, et al., and non-patent literature publication “*This is Watson*”, by D. A. Ferruci, et al., in the IBM Journal of Research and Development, Volume 56, Number 3/4, May/July 2012, are incorporated by reference in their entireties, including drawings.

BACKGROUND

Approximately 80% of online and digital data today is “unstructured data”, such as news articles, research reports, social media posts, and enterprise system data. Unlike “structured data”, e.g., databases, configuration tables, etc., which is readily useable by traditional computing processing, unstructured data is not directly compatible with traditional computer processes.

Understanding and interpreting unstructured data, such as electronic documents expressed in Natural Language (NL), is beyond the capacities of traditional search engines. Traditional search engines find keywords, and rank their findings according to the number of appearances of each keyword and their proximities to each other. In order to effectively use a keyword-based search engine, a user must input the most effective keywords. But, if the user does not know the correct keywords, the search engine may be of little use.

Further, keyword-based search engines have no ability to assign “right” or “wrong” to their results because they do not interpret their findings, and thus cannot detect disagreements between two or more search findings. For example, if a user is searching for a likely cause of a particular abdominal malady, he or she may input the symptoms (abdominal pain, nausea, etc.) as keywords into a keyword-based search engine. Two documents may be found by the search engine, each of which has similar quantities of appearances of the keywords (references to the symptoms) and thus are ranked similar to each other. However, the documents may depart radically from each other in their explanations of the potential cause (allergy, food poisoning, cancer, etc.) of the symptoms. The user must now try to make sense of these documents and determine which is correct, if either.

SUMMARY OF THE INVENTION

Curation of a corpus of a cognitive computing system is performed interactively by reporting on user interface device to a user a parse tree illustration of discrepancies and corresponding assigned confidence factors detected between at least a portion of a first document and a second or more documents in the corpus. Responsive to a user selection of an illustrated discrepancy in the parse tree, a drill-down dialog is prepared and displayed which shows at least a text

2

string for the portion of the first document and at least one conflicting text string from the second or more documents, and which provides at least one user-selectable administrative action option for handling the detected discrepancy. Responsive to receipt of user selection of an administrative action option, the computing system performs the action to handle the detected discrepancy.

BRIEF DESCRIPTION OF THE DRAWINGS

The figures presented herein, when considered in light of this description, form a complete disclosure of one or more embodiments of the invention, wherein like reference numbers in the figures represent similar or same elements or steps.

FIG. 1 illustrates a specialized document ingestion system which, in one embodiment, incorporates a computing core or customized logical circuit for processing logic which performs a logical process, such as by executing program instructions.

FIG. 2 sets forth one potential embodiment of the present invention relative to a cognitive computing system.

FIG. 3 shows a clustering structure of documents in a general manner.

FIG. 4 illustrates one particular example of clustered documents as shown in FIG. 3.

FIGS. 5 and 6 illustrate entities extracted from example text strings using natural language processing.

FIG. 7 depicts discrepancy detection and assignment of confidence factors to each discrepancy according to the nature and type of conflict between the text strings.

FIG. 8 provides an example user interface for a top-level view of any conflicts within a newly-ingested document relative to the documents already contained in the corpus.

FIG. 9 shows an example user interface for a drill-down view of a user-selected conflict in which a user can see the linked conflicting text and perform one of several reconciliation functions within the corpus.

DETAILED DESCRIPTION OF EMBODIMENT(S) OF THE INVENTION

Problems Recognized

The present inventors have realized that certain newer computing technologies may be combined, modified and operated in novel ways to improve upon the state of the art in search engines as described in the Background. By leveraging cognitive computing technology, inclusion of new electronic information sources in Natural Language can be improved by automatically interpreting the new document so that it can be compared to existing documents already within a domain of knowledge. Through this interpretation, a search computer system can detect disagreement between the new document and one or more documents already within the search systems purview.

Cognitive Computing and Deep Search Engines. Cognitive computing systems employ Natural Language Processing (NLP) to understand unstructured data, as well as machine learning processes which adapt and grow as additional unstructured language documents are ingested into a corpus of the cognitive computing system. An information corpus refers to a collection of unstructured data, such as a collection of digital or digitized literature, encyclopedias, dictionaries, thesauri, user’s manuals, press releases, legal texts, medical reference materials, news articles, academic text books, electronic messages, web logs (“blogs”), and social media postings.

These types of unstructured information items are usually expressed in various forms of natural language, some of it being more formally formatted (e.g., grammatically correct sentences and paragraphs, usage of proper terminology and definitions according to linguistic dictionaries, etc.), and some of it being more informally formatted (e.g., sentence fragments, use of time-dependent terminology such as slang, etc.).

International Business Machine Corporation's (IBM's) Watson™ platform is one such example of a cognitive computing system. Google's DeepMind™ and HP's Autonomy™ may also be viewed by some as a cognitive computing systems. Some subject-matter-specific cognitive computing systems have been proposed, such as DeepListening's automated personality analysis that operates on a short-message corpus for a single or group of messaging service users, and others such as TheySay, Lexalytics, etc.

Such systems typically employ combinations of NLP, artificial intelligence (AI) and machine learning (ML) to provide Deep Semantic Analysis of large corpuses of unstructured data. Many are provided as a platform, such as a networked computing service (e.g., on-demand computing, cloud computing, etc.) which can be engaged by other computer systems and users via an application programming interface (API).

Cognitive computing systems go far beyond typical keyword-based search engines by allowing users (human users and other computer processes) to ask complex questions. In response to a natural language query, a cognitive computing system analyzes the unstructured data in its corpus using NLP to understand grammar and context of each information item, and it presents candidate answers and/or solutions to the user ranked by certainty of correctness.

Before attempting to answer any questions for a user, a cognitive computing system must "learn" a new subject. It does so by loading or "ingesting" a plurality of unstructured language documents, electronic messages, web pages, etc., into the corpus, and question and answer pairs are added to train the system on the subject. As new unstructured language information becomes available, it may also be ingested into the corpus to expand the knowledgebase for future answers.

After an initial corpus is built and the system has been trained on the subject, when a user asks a new complex question of the system, it searches the corpus to find a plurality of potential answers. It also collects evidence within the corpus, such as how many sources agree on a particular possible answer, and rates the quality of the evidence according to a scoring process. Finally, potential answers which meet a threshold of confidence of being correct are ranked against each other and presented or returned to the user.

Referring to FIG. 2, one available embodiment of the present invention is integrated as a document ingestion pre-processor (250) of the corpus-building and textual resource acquisition function of a cognitive computing system, such as the functions described in the "Content Acquisition" section of "The AI Behind Watson—The Technical Article" by David A. Ferrucci, et al., (Association for the Advancement of Artificial Intelligence), and also described in "This is Watson", by Ferrucci, et al. (IBM Journal of Research and Development, Volume 56, Number 3/4, May/July 2012, see especially pages 4:1-4:11), both of which are incorporated by reference into the present disclosure. Other configurations and architectures may be produced by other cognitive computing systems, to which embodiments of the present invention may equally well provide improvement.

The exemplary improved cognitive computing system (200) is comprised of a computing system having one or more computer-readable memory devices, microprocessors, appropriate power supplies, computer network interfaces, and an operating system with application programs. One or more functions within the operating system and/or application programs implement machine logic functions (202-211).

Prior to handling input questions, corpus content (252) is acquired usually through a combination of manual and automatic steps. Example questions representative of the problem space (i.e., knowledge domain) are analyzed to produce a description of the kinds of questions that must be answered and a characterization of the application domain. Analyzing example questions may be manually performed, and domain analysis may be informed automatically by statistical analyses, such as the lexical answer type (LAT) analysis. For example, when IBM prepared its Watson cognitive computing system to play the Jeopardy game against human contestants, a corpus of information was prepared by ingesting a wide range of digital encyclopedias, dictionaries, thesauri, newswire articles, literary works, etc. (251).

After a baseline corpus (252) is established, an automatic corpus expansion process engages having four high-level steps:

- (1) identifying seed documents and retrieving related documents from the networked storage servers such as web servers;
- (2) extracting self-contained text fragments from the retrieved documents;
- (3) scoring the fragments based on whether they are informative relative to the original seed document; and
- (4) merging the most informative fragments into the information corpus.

During runtime, the cognitive computing system receives a question or query input (201), analyzes (202) the question and decompose (203) it. A primary search (204) on the corpus is performed, and candidate answers are correlated (205) and provided to a hypothesis generator (206). Optionally, evidence to support or discredit each hypothesis may be used to score (207) the candidate answers against the one or more candidate answers. Synthesis (208) and final merging and ranking (209) using trained models of the knowledge domain yield one or more outputs (210) which may be human readable (e.g., an answer to a question in a game show or a possible answer to a medical mystery, etc.), machine readable (e.g., via an API), or both.

Recognition of Benefits of Early Detection of Conflicting Information. The present inventors have realized that, during corpus initialization, some documents may be ingested which disagree with or are contradictory with some of the previously-ingested documents. Thus, they have discovered an opportunity to detect, flag and resolve these contradictions early in the corpus building process, thereby relieving and enhancing future search and classification steps. Some embodiments of the present invention will be realized as a document ingestion pre-processor (250) as shown in FIG. 2, although other embodiments may integrate the machine logic functions according to the invention in other points of a cognitive computing system architecture.

Exemplary embodiments of this invention relate to document ingestion (250) into a corpus for an information retrieval system, especially for cognitive computing systems such as but not limited to the IBM Watson™ computing system. In information retrieval systems, documents from which answers and information are retrieved form the cor-

pus. The accuracy of the information in the corpus is important, and the present inventors have recognized that it is desirable to have an ingestion tool that can automatically detect discrepancies between documents in a corpus. Embodiments according to the present invention include a discrepancy engine that can flag documents with discrepancies using a multi-phased process. A first phase is to map potential or previously asked questions to documents, sub-document, document sections, document fragments, and micro-documents, such as those identified by the National Institute of Standards and Technology Text Retrieval Conference Workshops referred to as TREC documents (TRECs) based on the source of an answer. In a second phase, the questions and corresponding documents are organized into categories using available document clustering techniques. In a third phase, when a new document is ingested into the existing corpus, it will be assigned to one or more of the existing categories, and fourth, the questions associated with that category are turned through the system again to find changes to the corpus caused by the addition of the new document.

Fifth, the corpus changes are used to identify documents with potential discrepancies among themselves, and sixth, an entities extraction process is used to generate individual statements for comparison by an ingestion engineer.

Example Usage and Operation. Multiple documents are ingested into a cognitive computing system to form a corpus as discussed in the previous paragraphs. Additional documents may be added to the corpus at a later time. During the ingestion process, with current technology (i.e., without an embodiment of the present invention), the user is not informed of any discrepancies that might exist among documents which have been ingested into the corpus. During runtime after corpus initialization, when an event is developing, or a long-term story is evolving in which causality and opinions differ, the present inventors have recognized the value of linking corpus documents or document sections based on whether they corroborate or disagree with each other and a hypothesis. Thus, the enhanced output of the cognitive computing system not only would include the candidate answers and their rankings, but also a set of linked documents which support each candidate answer and which controvert each candidate answer.

There are a variety of discrepancies that can be detected, such as material facts might differ (person/place/date), attribution of causality, and subjective factors such as sentiment, etc. The present inventors have realized that this could be valuable as a discovery tool to enrich the presentation of complex topics, such as presenting developing news events along a timeline, or in doing sociological analysis, or attributing the spin given to a story by different stakeholders, or in an enterprise scenario in identifying elements in conflict from various enterprise departments or across time. For example, in today's online news industry, there is a high degree of time pressure to get stories and reports published before the facts relating to them can be sufficiently validated. There might be a fact that comes out initially that is subsequently changed or discovered to be false. But, that change is interesting from analytical perspective such as characterizing the 'initial blame' versus final determination of the actual responsible party for a story or event.

In another example usage case, multiple versions of a product manual may include a lot of similar information, but may also present a small amount of different information. It can be difficult for human readers of such manuals to discern the differences between the two versions. For example, an Apple iPhone4™ user manual may be superseded by an

iPhone 5 user manual, but it is not clear due to the vast amount of similarity between the two what is actually changed or different. Using a cognitive computing system, one may ask "what is different between the iPhone 4 and 5 user manuals?", and the output results may show the user only the things that are nearly the same yet have some differing detail: such as a button moved or process is unexpected.

Still other workflows for an administrator managing the ingestion of new documents and document fragments into a cognitive computing system corpus is determining whether or not detected differences are really false or meant as humor, whether a particular source of information perhaps should be blocked from ingestion due to consistently incorrect or offensive information, and whether there are multiple versions of truths that could be promoted to an expert group for review or to an open discussion forum online such as a "wiki" group for further consideration and determination regarding which points are accurate.

First Example Embodiment: A Discrepancy Detection Engine The document ingestion pre-processor (250) in the present example embodiment would be added to a cognitive computing system as shown and discussed with respect to FIG. 2. Upon addition of a document or user-selection of an analyze option on a user interface, the ingested documents will be analyzed according to the logical processes and machine logic described herein, such as using a machine logic process shown in FIG. 1.

This particular embodiment is realized using a computing core (e.g., CPU 290, system memory 291, various communication and user interfaces 293, and one or more operating systems and optional embedded firmware 292) which performs a logical process (253-259), such as by executing program instructions, operating specialized integrated circuits, or a combination of executing instructions and operating circuits.

An ingested "document", for the purposes of the present disclosure, shall include one or more of each of the following: monolithic documents, large multi-section text documents, individual sections of a text, and micro-documents (e.g., tweets, instant messages, etc.).

The automated logical processes will typically include receiving a new document (big, little or fragment) (e.g., TREC documents) (251'), and then traversing all the previously-ingested documents (252) while searching for discrepancies amongst them relative to a knowledge domain. The discrepancies to be analyzed (255) may be identified via user-defined templates (254) for key relations, where a template might specify particular facts of interest in the usage-relevant knowledge domain(s) (253), or other criteria such as frequently cited/searched query terms or relationships. An example of such a template might be Person, Place, and Date.

The corpus documents (252) are then traversed for each of the supplied templates. Data from the documents will be supplied using the appropriate templates and saved for analysis. Once the data has been traversed and captured in template format, for example as extracted relationships or other annotations, it can be analyzed for any discrepancies. Sections of any documents that deviate from any other document will be highlighted in the User Interface or marked up in an electronic output to a non-human user (e.g., HTML or XML returned via an API). For each document that is highlighted as potentially different, the relevant documents will be identified and the section of the document will be identified.

If (256) no discrepancies are detected, the system may automatically ingest (257) the document, sub-document, document section, document fragment, or micro-document (251') into the corpus (252). However, found discrepancies will be flagged (251'') such as by highlighting the discrepant portions or passages, and displayed or presented to an administrator for approval (258) to ingest (250) the document with the discrepancies.

Optionally (not shown), the links between the discrepant documents may be recorded into a computer memory structure such as a database for future use in order to provide a user the agreeing and dissenting supporting documents in response to a question which elicits the linked documents.

The general machine logic process for discrepancy detection during ingestion of a new document, sub-document, document section, document fragment, or micro-document, such as a TREC format document, into a cognitive computing system corpus comprises several phases.

Phase 1. Populate a list of questions that are asked of the system, that are known to have both high confidence and highly rated answers, and for which the answers are in the previously ingested corpus.

- (a) Start with a ground truth collection from existing system which maps existing questions to known good answers aligned within the corpus that provides those answers.
- (b) Utilize other sources of known-good answers such as user feedback from running system, where the users have had a chance to review and validate the answers produced by the cognitive computing system.

The question-answering pipeline preserves (i) the questions previously asked, (ii) the answers returned in response to the previously-asked questions, (iii) the answer confidences associated with each answer, (iv) which answers were shown or otherwise provided to the user, (v) whether or not the evidence was viewed for each answer, (vi) the user feedback given for each answer, etc.

Phase 2. The system will align the answered questions with the document(s), sub-document(s), document section(s), document fragment(s), or micro-document(s) from which the correct answer came. A document in this case could be either an entire document that was ingested, or individual document sections, or individual tweets, or combinations thereof, for example. The cognitive computing system is capable of mapping the questions to a very fine level within the corpus at this point.

Phase 3. Organize the questions and document sections into clusters using a conventional clustering method such as K-means. Each cluster will be assigned a unique cluster ID, and clusters within clusters (sub-clusters) may also be created by the system.

For an example, refer now to FIG. 3 that shows a structure of a cluster of documents identified as "clust1". Within this cluster are three documents, shown as DocA, DocB and DocC. Further, DocB is associated with three document fragments, sections or micro-documents Trec1, Trec5 and Trec8. Similarly, DocC is associated with two fragments, sections or micro-documents Trec7 and Trec9. This structure (300) encapsulates the categorized questions as described in the previous paragraphs.

A more specific example is shown in FIG. 4 relative to the previously-discussed iPhone user's manual differences. The first cluster Clust1 include three questions: Do you sell iPhones? What model iPhones are available? Are iPhones the only mobile devices that Apple manufactures. Within the first cluster Clust 1 is a sub-cluster clust1_sub1 that includes

three questions: How can I charge my iPhone using a car charger? How do I charge my iPhone using a wall charger? How do I charge my iPhone using a USB cable? There may be also a second cluster clust2 of questions and a second sub-cluster, and so forth.

A categorized question may also be a child (member) of more than one cluster. So, for example, the question "How do I charge an iPhone using a wall charger" could be a child or member of a first cluster for "iPhone problems" and also be a child or member of a second cluster for "charging problems", etc. Once the questions have been mapped to clusters, then the system can assign labels to those clusters. In the example above, clust1 could be assigned a label "General iPhone questions", and sub-cluster clust1_sub1 could be assigned the label "iPhone charging methods".

At this point in the initialization of the system, the questions are mapped to clusters and documents mapped to clusters. Thus, the questions are associated to the related content of the documents in each cluster to which they belong.

Phase 4. After the corpus has been initialized in the first three phases, as new documents are ingested, the discrepancy detection steps are taken.

- (a) Categorize new documents. When a new document is ingested, it too will be categorized. Then, the questions which were previously associated with that topic or category are searched for conflicts with the content of the new document. Once the new document is ingested and turned into Trecs (or other sub-document unit types), and those Trecs are added to relevant clusters and optionally sub-clusters, the system applies questions associated with that cluster of the new Trec files. For example, assume Table 1 represents a simplified state of a corpus before a new document is ingested:

TABLE 1

Simplified State of a Corpus Before New Document Ingestion		
Cluster A	Cluster B	Cluster C
Doc A, Doc B, Doc C Q1, Q2, Q3	Doc D, Doc E, Doc F Q4, Q5	Doc G, Doc H, Doc I Q6, Q7

After ingestion of the new document, the corpus structure is amended as shown in Table 2:

TABLE 2

Simplified State of a Corpus After New Document Ingestion		
Cluster A	Cluster B	Cluster C
DocA, DocB, DocC, New Trec1 Q1, Q2, Q3	DocD, DocE, Doc F Q4, Q5	DocG, DocH, DocI, New Trec2 Q6, Q7

The Table 2 shows where the new Trec1 and Trec2 exist in the current state of the system for now. So, for new Trec1, the system queries the questions Q1, Q2 and Q3, while for the NewTrec2, the system queries the questions Q6 and Q7.

If the cognitive search finds that high confidence answer is returned for say Q1 when queried of NewTrec1, then the system declares that NewTrec1 is very closely related to any existing document that previously answered question Q1. The system can then look up which documents in ClusterA have correctly answered question Q1 (it may be more than one document).

So, for example, assume question Q1 has only been answered successfully by document DocA. The system can also ask negatives of the categorized questions, optionally using a method such as that disclosed in U.S. patent application Ser. No. 14/722,862, filed by Byron, et al, which is incorporated by reference into the present disclosure. High confidence answers for negative versions of the question would indicate that the document is possibly in conflict with the previously-ingested documents within that category or cluster. Perhaps in this example situation, the negative query of question Q3 returns a high confidence answer that Q3 is answered by document DocB.

(b) Identify documents for pairwise comparison. Given the foregoing example, the system now has identified two documents to investigate to look for possible conflicts because one document matched the new document using a positive question and another document matched the new document using a negated question from the cluster. The system will now compare the new document NewTrec1 with DocA and DocB to resolve this discrepancy, if it exists.

Turning to an example of how the system compares the new Trec1 with document DocA, assume that DocA, perhaps an electronic medical record, contains the text (501) of FIG. 5, and further assume that the new Trec1, perhaps an obituary listing from an electronic newspaper, contains the text (503). The system can perform entity extraction (502, 504) on the text (of both documents new Trec1 and DocA to yield entities person, dates, locations and causes (i.e. disease). However, the date “today” in the new Trec1 is relative, not absolute.

If no digital date, such as a date contained in metadata for Trec1 is available, the system can presume that “today” corresponds to “March 26th” because of the high degree of matches between the other extracted entities.

If, however, another date, such as a copyright date or a metadata data associated with new Trec1 or another statement found in another document already in the corpus may controvert the date of March 26th, such as the statement (601) with the entities (602) as shown in FIG. 6. In such a case, the system detects and flags this discrepancy.

Similarly, relationship extraction may be performed to establish natural language relationships between extracted entities, and discrepancies may be detected and flagged according to relationship disagreements between newly ingested documents and documents already in the corpus.

For example, consider two relationships parse tree structures Eq. 1 and Eq. 2 which may be generated from the foregoing text examples (501, 503, respectively) using natural language processing which was produced using the openly-available Link Grammar Parser by Davy Temperley, Daniel Sleator and John Lafferty, of Carnegie Mellon University, which is a syntactic parser of English based on “link grammar”. Given a sentence, the Link Grammar Parser system assigns to it a syntactic structure, which consists of a set of labeled links connecting pairs of words, producing a constituent (noun phrases, verb phrases, etc.) representation of a sentence:

(S (NP John Doe)	Eq. 1
(VP died March 26th	
(PP at	
(NP UMass Medical Center))	
(PP from	
(NP Stage 4 Cancer))))	
(S (NP John Doe)	Eq. 2
(VP , leaving	

-continued

(NP a wife a child) ,))
 (VP died March 27th
 (PP at (NP UMass))))

From these extracted relationships using a parser such as this, for example, the embodiment of the present invention can establish that subject entity and the date entities are related by the same action entity values (i.e., died), so the extracted relationships in these examples can be simplified as shown in FIG. 7. In the first pairwise comparison of the dates (701, 702) of the action, a high degree of confidence of a conflict or discrepancy can be assigned by the system (705) and reported.

However, in the second pairwise comparison of the location entities (703, 704), there is some similarity between the text strings (e.g., they both contain the sub-string “UMass”), so the discrepancy may be flagged and given a lower confidence of being a meaningful conflict (706).

Second Example Embodiment: A Discrepancy User Interface. The foregoing phases of a discrepancy detection engine embodiment may be combined with a user interface (UI) embodiment as follows. Those ordinarily skilled in the arts, however, will recognize that the engine and the UI embodiments may be realized and found useful separately, as well, wherein the unifying inventive concept relates to the detection, flagging and confidence rating advantages of the present invention.

Modern word processor application programs allow a user to compare two versions of the same document, and the user interface illustrates to a user the most recent changes. For example, if a phrase is changed in a later version of a document, the user interface of the word processor may show the previous version of the phrase with special marking such as strikethrough text, and may show the next version of the phrase with special marking such as underlining. However, the comparison methods are greatly dependent on the two documents being mostly the same with minor differences. Comparing two documents which are structurally different, such as an obituary and a medical record, would result in one document being entirely shown as deleted (strikeout text) and the other document’s text as being inserted (underlining). This is because the word processor application program’s comparison routine compares words and strings literally to each other without regard or understanding of the semantic or natural language meaning of the information contained within the compared documents.

Thus, a user interface according to the present invention is not well-served by the user interface model of word processors. The present inventors, therefore, hereby disclose an inventive user interface such as the example (800) illustrated in FIG. 8 for display on a computer display. A generic computer display does not normally provide such a user interface as described in the following paragraphs, so the available embodiments of the present invention may be realized through programming enhancements which are executed by a computer processor at the application level or the operating system level, or even embedded within a graphics display accelerator device, integrated circuit or chip.

The dialog box (804) is preferably displayed on a portion of a computer display in a manner that conveys the topically-clustered documents, document portions, and micro-documents as described in the foregoing paragraphs. Color may be used to visually set apart conflicts (i.e. red), similarities

11

with differences (i.e. yellow) and agreements (i.e. green) between a newly-ingested document and the documents already in the system's corpus. Still further, the flagging declaration text (801, 802, 803) may be selectable or clickable by a user, such as by manipulation of keys, buttons and/or pointing device(s), to engage a deeper investigation of the discrepancies and similarities which were detected during ingestion processing.

For example, such as in the engine embodiments disclosed in the foregoing paragraphs in which links between agreeing and controverting documents in the corpus were recorded during ingestion processing, clicking on or selecting the text (801) declaring a conflict between existing DocA in the corpus and the newly-ingested Trec1 may lead to another user information display on the computer display as shown (900) in FIG. 9. A generic computer display does not normally provide such a user interface as described in the following paragraphs, so the available embodiments of the present invention may be realized through programming enhancements which are executed by a computer processor at the application level or the operating system level, or even embedded within a graphics display accelerator device, integrated circuit or chip.

In this drill-down dialog (904), two or more conflicting elements, relationships, or a combination of elements and relationships may be shown (901, 903), along with one or more of the confidence ratings (901'), and one or more commands (902, 905) for user operation to reconcile the differences within the corpus. For example, the foregoing date of death discrepancy may be removed from the newly-ingested Trec1 document, replaced with the date value already found in the corpus, edited manually, ignored, ignored with specially notation, etc.

Using such a top-level summary user interface (800) with one or more drill-down user interfaces (900) for each of the items in the cluster model, an administrator or information engineer may selectively inspect and optionally resolve discrepancies in order to prepare the corpus in advance of querying it by a cognitive computing system.

Optional Tracking of Resolutions. In some embodiments, such as for applications in which information traceability is required or desired, each change or correction made via the user interfaces to elements and relationships detected as discrepancies can be tracked and recorded into a data structure in a computer memory. In implementation, such change tracking functions may resemble the "undo" list commonly found in word processor application programs and in database monitoring systems.

Computer Program Product Embodiments. The present invention may be a system, a method, and/or a computer program product at any possible technical detail level of integration. The computer program product may include a computer readable storage medium (or media) having computer readable program instructions thereon for causing a processor to carry out aspects of the present invention.

The computer readable storage medium can be a tangible device that can retain and store instructions for use by an instruction execution device. The computer readable storage medium may be, for example, but is not limited to, an electronic storage device, a magnetic storage device, an optical storage device, an electromagnetic storage device, a semiconductor storage device, or any suitable combination of the foregoing. A non-exhaustive list of more specific examples of the computer readable storage medium includes the following: a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory

12

(EPROM or Flash memory), a static random access memory (SRAM), a portable compact disc read-only memory (CD-ROM), a digital versatile disk (DVD), a memory stick, a floppy disk, a mechanically encoded device such as punch-cards or raised structures in a groove having instructions recorded thereon, and any suitable combination of the foregoing. A computer readable storage medium, as used herein, is not to be construed as being transitory signals per se, such as radio waves or other freely propagating electromagnetic waves, electromagnetic waves propagating through a waveguide or other transmission media (e.g., light pulses passing through a fiber-optic cable), or electrical signals transmitted through a wire.

Computer readable program instructions described herein can be downloaded to respective computing/processing devices from a computer readable storage medium or to an external computer or external storage device via a network, for example, the Internet, a local area network, a wide area network and/or a wireless network. The network may comprise copper transmission cables, optical transmission fibers, wireless transmission, routers, firewalls, switches, gateway computers and/or edge servers. A network adapter card or network interface in each computing/processing device receives computer readable program instructions from the network and forwards the computer readable program instructions for storage in a computer readable storage medium within the respective computing/processing device.

Computer readable program instructions for carrying out operations of the present invention may be assembler instructions, instruction-set-architecture (ISA) instructions, machine instructions, machine dependent instructions, microcode, firmware instructions, state-setting data, configuration data for integrated circuitry, or either source code or object code written in any combination of one or more programming languages, including an object oriented programming language such as Smalltalk, C++, or the like, and procedural programming languages, such as the "C" programming language or similar programming languages. The computer readable program instructions may execute entirely on the user's computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user's computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider). In some embodiments, electronic circuitry including, for example, programmable logic circuitry, field-programmable gate arrays (FPGA), or programmable logic arrays (PLA) may execute the computer readable program instructions by utilizing state information of the computer readable program instructions to personalize the electronic circuitry, in order to perform aspects of the present invention.

Aspects of the present invention are described herein with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems), and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer readable program instructions.

These computer readable program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the

13

instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks. These computer readable program instructions may also be stored in a computer readable storage medium that can direct a computer, a programmable data processing apparatus, and/or other devices to function in a particular manner, such that the computer readable storage medium having instructions stored therein comprises an article of manufacture including instructions which implement aspects of the function/act specified in the flowchart and/or block diagram block or blocks.

The computer readable program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other device to cause a series of operational steps to be performed on the computer, other programmable apparatus or other device to produce a computer implemented process, such that the instructions which execute on the computer, other programmable apparatus, or other device implement the functions/acts specified in the flowchart and/or block diagram block or blocks.

The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods, and computer program products according to various embodiments of the present invention. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of instructions, which comprises one or more executable instructions for implementing the specified logical function(s). In some alternative implementations, the functions noted in the blocks may occur out of the order noted in the Figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts or carry out combinations of special purpose hardware and computer instructions.

Conclusion. The terminology used herein is for the purpose of describing particular embodiments only and is not intended to be limiting of the invention. As used herein, the singular forms “a”, “an” and “the” are intended to include the plural forms as well, unless the context clearly indicates otherwise. It will be further understood that the terms “comprises” and/or “comprising,” when used in this specification, specify the presence of stated features, steps, operations, elements, and/or components, but do not preclude the presence or addition of one or more other features, steps, operations, elements, components, and/or groups thereof, unless specifically stated otherwise.

The corresponding structures, materials, acts, and equivalents of all means or step plus function elements in the claims below are intended to include any structure, material, or act for performing the function in combination with other claimed elements as specifically claimed. The description of the present invention has been presented for purposes of illustration and description, but is not intended to be exhaustive or limited to the invention in the form disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art without departing from the scope and spirit of the invention. The embodiment was chosen and described in order to best explain the principles of the

14

invention and the practical application, and to enable others of ordinary skill in the art to understand the invention for various embodiments with various modifications as are suited to the particular use contemplated.

It should also be recognized by those skilled in the art that certain embodiments utilizing a microprocessor executing a logical process may also be realized through customized electronic circuitry performing the same logical process or processes.

It will be readily recognized by those skilled in the art that the foregoing example embodiments do not define the extent or scope of the present invention, but instead are provided as illustrations of how to make and use at least one embodiment of the invention. The following claims define the extent and scope of at least one invention disclosed herein.

What is claimed is:

1. A method for curating an information corpus of a cognitive computing system, the method comprising the steps of:

producing on a user interface device of a computing system a parse tree illustration of one or more discrepancies and confidence factors previously detected between at least a portion of a first document and a second or more documents in the information corpus, wherein the portion and the second or more documents comprise unstructured information items, wherein the parse tree illustration comprises a structure of a cluster of documents encapsulating categorized previously-asked questions about the information corpus, and

wherein the confidence factors are previously-assigned corresponding to the questions and the documents; responsive to receipt, by a computing system, from the user interface device, of a selection of an illustrated discrepancy in the parse tree, displaying on the user interface device, a drill-down dialog which shows at least one of the previously-detected discrepancies; and responsive to receipt, by a computing system, from the user interface device, of an administrative action option, performing by the computing system the administrative action to curate the information corpus.

2. The method as set forth in claim 1 further comprising, responsive to the handling of the discrepancy, recording by a computing system, the administrative action performed so that the discrepancy may optionally be suppressed from future display or reviewed on the user interface device in the future.

3. The method as set forth in claim 1 wherein the displaying of the drill-down dialog comprises accessing one or more recorded links between the first and the second or more documents and creating a graphical user display according to the links.

4. The method as set forth in claim 1 wherein the discrepancies reflect conflicting answers provided by the first and the second or more documents when queried using a question associated with a category to which the first and second or more documents are assigned.

5. The method as set forth in claim 1 wherein the administrative action option comprises one or more actions selected from the group consisting of flagging the first document, removing the first document, replacing the first document, editing the first document, ignoring the discrepancy, flagging the second or more document, removing the second or more document, replacing the second or more document, and editing the second or more document.

15

6. The method as set forth in claim 1 wherein at least one of the first, second or more documents comprise a fragment of a document.

7. A computer program product for a user to curate an information corpus of a cognitive computing system comprising:

a computer-readable memory device which is not a propagating signal per se; and
program instructions embodied by the computer-readable memory device which, when executed, cause a processor to perform steps of:

produce, on a user interface device of a computing system a parse tree illustration of one or more discrepancies and confidence factors previously detected between at least a portion of a first document and a second or more documents in the information corpus,

wherein the portion and the second or more documents comprise unstructured information items,

wherein the parse tree illustration comprises a structure of a cluster of documents encapsulating categorized previously-asked questions about the information corpus, and

wherein the confidence factors are previously-assigned corresponding to the questions and the documents; responsive to receipt, from the user interface device, of a selection of an illustrated discrepancy in the parse tree, display, on the user interface device a drill-down dialog which shows at least one of the one or more discrepancies; and

responsive to receipt, from the user interface device, of an administrative action option, perform the administrative action to curate the corpus.

8. The computer program product as set forth in claim 7 wherein the program instructions further comprise program instructions for, responsive to the handling of the discrepancy, recording the administrative action performed so that the discrepancy may optionally be suppressed from future display or reviewed on the user interface device in the future.

9. The computer program product as set forth in claim 7 wherein the program instructions for displaying of the drill-down dialog comprises program instructions for accessing one or more recorded links between the first and the second or more documents and creating a graphical user display according to the links.

10. The computer program product as set forth in claim 7 wherein the discrepancies reflect conflicting answers provided by the first and the second or more documents when queried using a question associated with a category to which the first and second or more documents are assigned.

11. The computer program product as set forth in claim 7 wherein the administrative action option comprises one or more actions selected from the group consisting of flagging the first document, removing the first document, replacing the first document, editing the first document, ignoring the discrepancy, flagging the second or more document, removing the second or more document, replacing the second or more document, and editing the second or more document.

12. The computer program product as set forth in claim 7 wherein at least one of the first, second or more documents comprise a fragment of a document.

16

13. A system for a user to curate an information corpus of a cognitive computing system comprising:

a computing system having a processor;

a computer-readable memory device which is not a propagating signal per se; and

program instructions embodied by the computer-readable memory device which, when executed, cause the processor to perform steps of:

produce, on a user interface device of a computing system, a parse tree illustration of one or more discrepancies and confidence factors previously detected between at least a portion of a first document and a second or more documents in the information corpus,

wherein the portion and the second or more documents comprise unstructured information items, wherein the parse tree illustration comprises a structure of a cluster of documents encapsulating categorized previously-asked questions about the information corpus, and

wherein the confidence factors are previously-assigned corresponding to the questions and the documents; responsive to receipt, from the user interface device, of a selection of an illustrated discrepancy in the parse tree, display on the user interface device, a drill-down dialog which shows at least one of the one or more discrepancies; and

responsive to receipt, from the user interface device, an administrative action option, perform the administrative action to curate the information corpus.

14. The system as set forth in claim 13 wherein the program instructions further comprise program instructions for, responsive to the handling of the discrepancy, recording the administrative action performed so that the discrepancy may optionally be suppressed from future display or reviewed on the user interface device in the future.

15. The system as set forth in claim 13 wherein the program instructions for displaying of the drill-down dialog comprises program instructions for accessing one or more recorded links between the first and the second or more documents and creating a graphical user display according to the links.

16. The system as set forth in claim 13 wherein the discrepancies reflect conflicting answers provided by the first and the second or more documents when queried using a question associated with a category to which the first and second or more documents are assigned.

17. The system as set forth in claim 13 wherein the administrative action option comprises one or more actions selected from the group consisting of flagging the first document, removing the first document, replacing the first document, editing the first document, ignoring the discrepancy, flagging the second or more document, removing the second or more document, replacing the second or more document, and editing the second or more document.

18. The system as set forth in claim 13 wherein at least one of the first, second or more documents comprise a fragment of a document.

* * * * *